

Storage for High-Performance Computing Gets Enterprise Ready

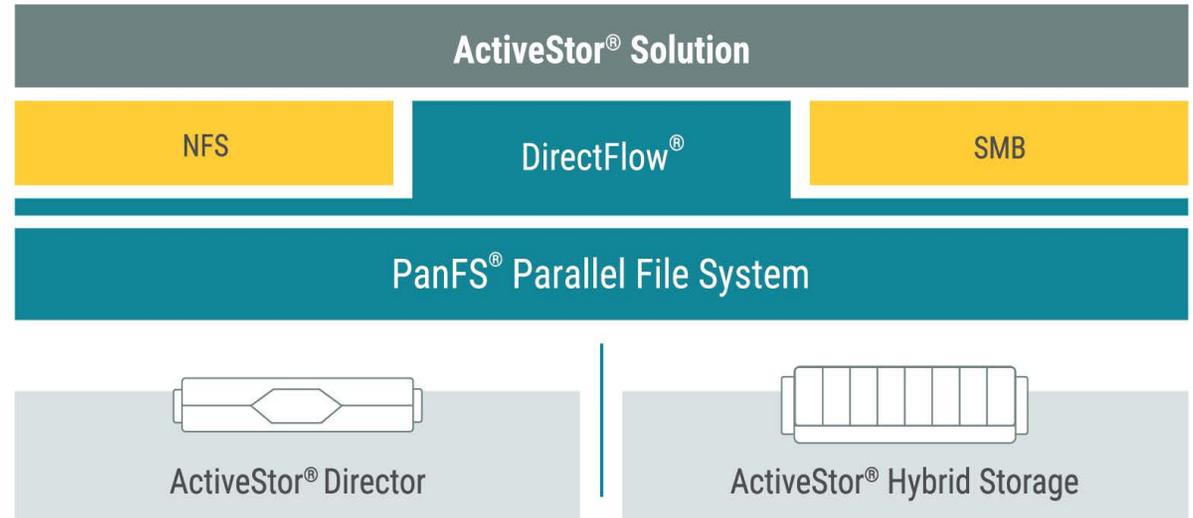
by George Crump

High-Performance Computing (HPC) is the proving grounds to test a storage system's ability to handle mixed workloads. While HPC is an important market it is relatively small, at least in comparison to the broader enterprise data center storage market. With initiatives like Big Data and Internet of Things (IoT) though, more enterprises are beginning to stand up architectures that look very similar to traditional HPC designs.

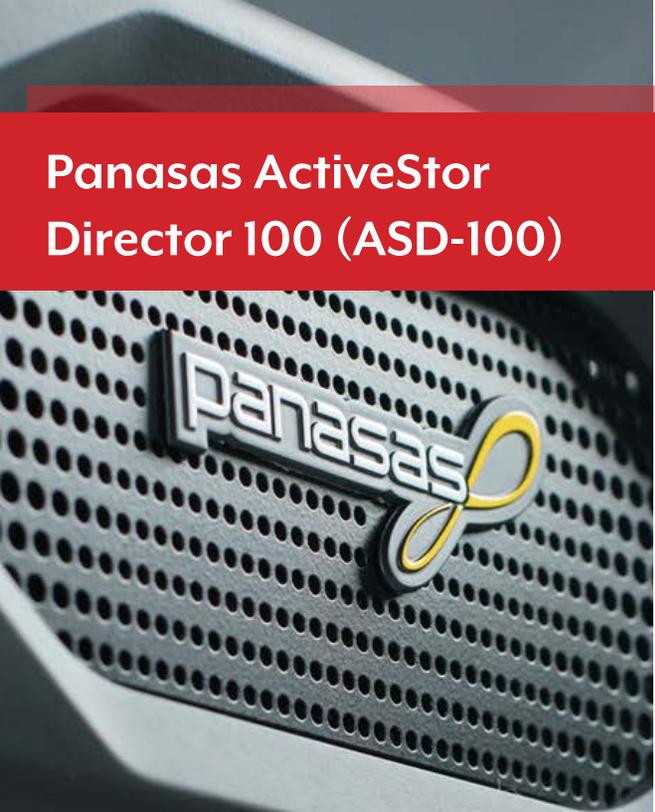
The Traits of HPC Storage

HPC Storage Systems have to respond to a wide range of workload types. Requirements for both high transaction file IO and high bandwidth are not uncommon. These environments often have massive metadata requirements with very high file counts and very elaborate directory structures.

For the users of HPC storage, it all comes down to productivity. How many more simulations or models can users run in the same amount of time versus another solution? Parallel access to data is a key requirement so data response time is consistent no matter how many users are accessing the system.



Given these attributes, the HPC storage system should be more than just scale-up. It should be parallel in nature, while still supporting legacy access via NFS or SMB. Also, these systems should have the option to manage metadata separately from actual data, so the number of files or the sophistication of the directory structure does not impact response time.



Panadas ActiveStor Director 100 (ASD-100)

Panadas is a veteran of the HPC storage market and remains tightly focused on it. Its current generation of hardware was scale out in design, but each shelf had a dedicated director that managed metadata and cluster services. Now, Panadas is introducing the ASD-100 which supercharges the brain power of the Panadas parallel storage system.

The ASD-100 is now disaggregated from the traditional nodes to unlock its full performance potential. This disaggregation also provides flexibility to support evolving network infrastructure requirements. The ASD-100 is a 2U chassis with four nodes inside and is still delivered on industry standard hardware. It can co-exist with the prior generation of Panadas hardware so upgrades are seamless.



Panadas ActiveStor Hybrid 100 (ASH-100)

The ASH-100 is now primarily focused on providing storage capacity to the environment, managed by the director. Removal of the per-shelf metadata processor frees up space for additional capacity. Also the ASH-100 can be configured with two 12TB HDDs and a 1.9TB SSD in a hybrid configuration. Each blade includes an Intel Atom processor with 16GB of DRAM room for two HDDs and 1 SSD.

The Panadas software uses the SSD as more than just a cache, in that it specifically places small files on the SSD and large files on the HDD. Considering that most large files are sequentially read and most small files are randomly read this configuration should provide optimal performance per-blade.

PanFS 7.0

Also updated is Panasas' file system, PanFS 7.0, which is the parallel file system at the heart of the solution. It brings PanFS' core operating system, FreeBSD, to the current level. Doing so means an improved NFS server implementation and a simpler, more dynamic GUI.

Panasas has also updated its DirectFlow Protocol with improved memory allocation and an improved read-ahead feature. The result is a 15% improvement in performance.

STORAGESWISS TAKE

Panasas is at the intersection of high-performance computing and scale-out NAS. Its timing is fortuitous as the HPC market continues to grow by about 7% with most of that growth coming from enterprises. But those enterprises also want solutions that don't have the lab experiment feel of many HPC solutions. Panasas gives them an enterprise experience with HPC capabilities.

THE ANALYST

George Crump is the founder of Storage Switzerland, the leading storage analyst focused on the subjects of big data, solid state storage, virtualization, cloud computing and data protection. He is widely recognized for his articles, white papers, and videos on such current approaches as all-flash arrays, deduplication, SSDs, software-defined storage, backup appliances, and storage networking. He has over 25 years of experience designing storage solutions for data centers across the U.S.



THE FIRM



Storage Switzerland is the leading storage analyst firm focused on the emerging storage categories of memory-based storage (Flash), Big Data, virtualization, and cloud computing. The firm is widely recognized for its blogs, white papers and videos on current approaches such as all-flash arrays, deduplication, SSD's, software-defined storage, backup appliances and storage networking. The name "Storage Switzerland" indicates a pledge to provide neutral analysis of the storage marketplace, rather than focusing on a single vendor approach.

ABOUT OUR PARTNER



Panasas is a premier provider of high-performance storage solutions. Their ActiveStor® scale-out network-attached storage (NAS) supports industry and research innovation around the world, with the fastest plug-and-play parallel data storage system, optimized to accelerate workflows, simplify data management, and deploy easily as an appliance. Delivered as a fully integrated clustered appliance solution, ActiveStor incorporates flash and SATA storage nodes, a distributed performance-optimized file system, and client protocols. Using the Panasas PanFS® parallel file system and Panasas DirectFlow® parallel data access protocol, ActiveStor redefines what it means to scale. Performance, data protection, and manageability all grow as the solution scales.